

Explanation in Recommender Systems

DAVID MCSHERRY

*School of Computing and Information Engineering, University of Ulster, Coleraine
BT52 1SA, Northern Ireland, UK (e-mail: dm.mcsherry@ulster.ac.uk)*

Abstract. There is increasing awareness in recommender systems research of the need to make the recommendation process more transparent to users. Following a brief review of existing approaches to explanation in recommender systems, we focus in this paper on a case-based reasoning (CBR) approach to product recommendation that offers important benefits in terms of the ease with which the recommendation process can be explained and the system's recommendations can be justified. For example, recommendations based on incomplete queries can be justified on the grounds that the user's preferences with respect to attributes not mentioned in her query cannot affect the outcome. We also show how the relevance of any question the user is asked can be explained in terms of its ability to discriminate between competing cases, thus giving users a unique insight into the recommendation process.

Keywords: attribute-selection strategy, case-based reasoning, explanation, recommender systems

1. Introduction

The importance of intelligent systems having the ability to explain their reasoning is well recognised in domains such as medical decision making and intelligent tutoring (e.g. Armengol et al., 2001; Sørmo and Aamodt, 2002; Evans-Romaine and Marling, 2003). In an intelligent tutoring system, for example, communicating the reasoning process to students may be as important as finding the right solution. Until recently, explanation in recommender systems appears to have been a relatively neglected issue. However, recent research has highlighted the importance of making the recommendation process more transparent to users and the potential role of explanation in achieving this objective (Herlocker et al., 2000; Shimazu, 2002; McSherry, 2002b, 2003b; Reilly et al., 2005).

Herlocker et al. (2000) suggest that the *black box* image of recommender systems may be one of the reasons why they have gained much less acceptance in high-risk domains such as holiday packages or investment portfolios than in low-risk domains such as CDs or movies. They

argue that extracting meaningful explanations from the computational models on which recommendations are based is a challenge that must be addressed to enable the development of recommender systems that are more understandable, more effective, and more acceptable. It is an argument that seems equally compelling in collaborative and content-based approaches to product recommendation.

McSherry (2003a) proposes a case-base reasoning (CBR) approach to product recommendation that combines an effective strategy for reducing the length of recommendation dialogues with a mechanism for ensuring that the dialogue is terminated only when it is certain that the recommendation will be the same no matter how the user chooses to extend her query. Referring to the approach as *incremental* nearest neighbour (iNN), we focus here on the benefits it offers in terms of making the recommendation process more transparent to users. One advantage is that recommendations based on incomplete queries can be justified on the grounds that the user's preferences with respect to attributes not mentioned in her query cannot affect the outcome. We also show how the relevance of any question the user is asked can be explained in terms of its ability to discriminate between competing cases, thus giving users a unique insight into the recommendation process.

In Section 2, we examine existing approaches to explanation in recommender systems and some of the lessons learned from this research. In Section 3, we present a detailed account of the recommendation process in iNN and the important role played by the concept of *case dominance* in the approach. In Section 4, we present an approach to explanation in which there is no requirement for domain knowledge other than the similarity knowledge and cases already available to the system. We demonstrate the approach in a mixed-initiative recommender system called *Top Case* which can explain the relevance of any question the user is asked in strategic terms, recognise when the dialogue can be safely terminated, and justify its recommendations on the grounds that any un-elicited preferences of the user cannot affect the outcome. Related work is discussed in Section 5 and our conclusions are presented in Section 6.

2. Existing Approaches

Herlocker et al. (2000) evaluated several approaches to explanation in the collaborative movie recommender *MovieLens* in terms of their

effects on user acceptance of the system's recommendations. The most convincing explanation of why a movie was recommended was one in which users were shown a histogram of the ratings of the same movie by similar users. Moreover, grouping together of good ratings (4 or 5) and bad ratings (1 or 2) and separation of ambivalent ratings (3) was found to increase the effectiveness of the histogram approach. Interestingly, the second most convincing explanation was a simple statement of the system's performance in the past e.g.

MovieLens has predicted correctly for you 80% of the time in the past.

Another important finding was that some of the explanations evaluated had a *negative* impact on acceptance, the goal of explanation in this instance, showing that no explanation may be better than one that is poorly designed.

CBR recommender systems that can explain their recommendations include Shimazu's (2002) *ExpertClerk* and McSherry's (2003b) *First Case*. *ExpertClerk* can explain why it is proposing two contrasting products in terms of the trade-offs between their positive and negative features e.g.

This blouse is more expensive but the material is silk. That one is cheaper but the material is polyester.

Its explanations are based on assumed preferences with respect to attributes not mentioned in the user's query. For example, a blouse made of silk is assumed to be preferred to one made of polyester.

In a similar way, *First Case* can explain why one case is more highly recommended than another by highlighting the benefits it offers (McSherry, 2003b). As the following example illustrates, it can also explain why a given product, such as a personal computer, is recommended in terms of the *compromises* it involves with respect to the user's preferences.

Case 38 differs from your query only in processor speed and monitor size. It is better than Case 50 in terms of memory and price.

However, the potential role of explanation in recommender systems is not limited to explaining why a particular item is recommended. In this paper, we present a CBR recommender system that can also explain the relevance of any question the user is asked in terms of its ability to discriminate between competing cases. Reilly et al.'s (2005) dynamic approach to critiquing in recommender systems differs from traditional critiquing approaches (e.g. Burke, 2002)

in that critiques are dynamically generated by the system and may involve compromises as well as improvements relative to the currently recommended case. In this way, the user is informed in advance of trade-offs associated with desired improvements. Before selecting a suggested critique, the user can ask to see an explanation of the trade-offs involved.

In recommender systems that treat some or all of the user's requirements as constraints that must be satisfied, explanation can also play an important role in recovery from the retrieval failures that occur when there is no exact match for the user's requirements (Hammond et al., 1996; McSherry, 2004). Hammond et al.'s (1996) *Car Navigator* is a recommender system for cars that uses declarative knowledge to explain trade-offs that are known to be common causes of retrieval failure in the domain, such as that between fuel economy and horsepower. For example, if the user asks for good fuel economy and high horsepower, she is shown a movie explaining the trade-off between these features. The user is also advised that she will need to revise her query if she hopes to find a car that meets her requirements.

In recent work, we combined a *knowledge-light* approach to explanation of retrieval failure with a mixed-initiative approach to recovery from retrieval failure in a CBR recommender system called *ShowMe* (McSherry, 2004). Failure to retrieve a case that exactly matches the user's query triggers an explanation that draws the user's attention to combinations of features in her query for which there are no matching cases e.g.

Sorry, there are no products that match these combinations of features in your query: (price \leq 500, type = laptop), (type = laptop, screen size = 19)

As well as highlighting areas of the product space in which the case library is lacking in coverage, the explanation may reveal *misconceptions* on the part of the user such as the price she expects to pay for the product she is seeking. Showing the user only the minimally failing *sub-queries* of her query, a technique we have adapted from research on co-operative responses to failing database queries (Gasterland et al., 1992), helps to minimise cognitive load in the approach. Explanation of the retrieval failure is followed in *ShowMe* by a mixed-initiative recovery process in which the user is guided in the selection of one or more constraints to be eliminated from her query (McSherry, 2004).

3. Incremental Nearest Neighbour

In this section, a brief overview of conversational CBR (Aha et al., 2001) in the context of product recommendation is followed by a detailed account of the recommendation process in iNN and the important role played by the concept of *case dominance* in the approach. One distinguishing feature of our approach is a *goal-driven* attribute selection strategy that has been shown to be very effective in reducing the length of recommendation dialogues (McSherry, 2003a). Another is a simple mechanism for ensuring that the dialogue is terminated only when it is certain that a more similar case will not be found if the dialogue is allowed to continue.

3.1. Conversational CBR

In CBR approaches to product recommendation, descriptions of the available products are stored in a case library and retrieved in response to a query representing the preferences of the user. In *conversational CBR* (CCBR) approaches like iNN, a query is incrementally (and often incompletely) elicited in an interactive dialogue with the user. We focus here on approaches in which the retrieval of recommended cases is based on their *similarity* to the elicited query, rather than relying on exact matching as in most decision-tree approaches (e.g. Doyle and Cunningham, 2000; McSherry, 2001b).

Given a query Q over a subset A_Q of the case attributes A , the similarity of any case C to Q is typically defined to be:

$$\text{sim}(C, Q) = \sum_{a \in A_Q} w_a \text{sim}_a(C, Q),$$

where for each $a \in A$, w_a is the importance weight assigned to a and $\text{sim}_a(C, Q)$ is a *local* measure of the similarity of $\pi_a(C)$, the value of a in C , to $\pi_a(Q)$, the preferred value of a . As usual in practice, we assume that $0 \leq \text{sim}_a(x, y) \leq 1$ for all $a \in A$ and that $\text{sim}_a(x, y) = 1$ if and only if $x = y$. We also assume that for each $a \in A$, the distance measure $1 - \text{sim}_a$ satisfies the *triangle* inequality.

A generic algorithm for CCBP in product recommendation (CCBR-PR) is shown in Figure 1. At each stage of the recommendation dialogue, the system selects the next most useful attribute, asks the user for the preferred value, and retrieves the case (or product)

```

algorithm CCBP-PR
begin
   $Q \leftarrow \emptyset$ 
  repeat
    select the next most useful attribute  $a$ 
    ask the user for the preferred value  $v$  of  $a$ 
     $Q \leftarrow Q \cup \{a = v\}$ 
    retrieve the case  $C$  that is most similar to  $Q$ 
    recommend  $C$ 
  until termination criteria are satisfied
end

```

Figure 1. Conversational CBR in product recommendation.

that is most similar to the query that has been elicited so far. The dialogue continues until some predefined *termination criteria* are satisfied, or until no further attributes remain. The case recommended on each cycle is usually the one that is most similar to the current query. However, it is not unusual for more than one case to be maximally similar to a given query, in which case we assume that all such cases are equally recommended. That is, we define the *recommendation* for a given query Q to be:

$$r(Q) = \{C : \text{sim}(C, Q) \geq \text{sim}(C^\circ, Q) \text{ for all } C^\circ\}.$$

Cases other than those that are maximally similar to the current query may also be presented as alternatives that the user may wish to consider, though the number of cases that can be presented to the user may be limited by the available screen space. Of course, cognitive load is another important consideration.

The defining components of a CCBP-PR algorithm are the strategy used to select the most useful attribute on each recommendation cycle and the criteria used to decide when the dialogue should be terminated. Possible approaches to attribute selection include giving priority to the most important of the remaining attributes (McSherry, 2003a) and the similarity-based approach proposed by Kohlmaier et al. (2001). Various approaches to termination of the recommendation dialogue are also possible. For example, the dialogue could be terminated when the current query Q is such that $|r(Q)| = 1$ or when the similarity of any case reaches a predefined threshold. As we shall see in Section 3.4, the criteria for termination of the

recommendation dialogue in iNN are closely linked to the attribute-selection strategy that characterises the approach.

3.2. Identifying dominated cases

In iNN, an important role in the recommendation process is played by the concept of case dominance that we now define.

Definition 1: *A given case C_2 is dominated by another case C_1 with respect to a query Q if $\text{sim}(C_2, Q) < \text{sim}(C_1, Q)$ and $\text{sim}(C_2, Q^*) < \text{sim}(C_1, Q^*)$ for all extensions Q^* of Q .*

One reason for the importance of case dominance in product recommendation is that if a given case C_2 is dominated by another case C_1 then the product represented by C_2 can be eliminated. Of course, the number of ways in which a given query can be extended may be very large. So given an incomplete query Q and cases C_1, C_2 such that $\text{sim}(C_2, Q) < \text{sim}(C_1, Q)$, how can we tell if C_2 is dominated by C_1 without resorting to exhaustive search?

One situation in which C_2 is clearly dominated by C_1 is when both cases have the same values for all the remaining attributes. Another is when $\text{sim}(C_1, Q) - \text{sim}(C_2, Q)$ is greater than the sum of the importance weights of all the remaining attributes. In situations where dominance is less obvious, account must be taken of the similarity between the two cases as well as their similarities to the current query (McSherry, 2003a). The criterion used to identify dominated cases in iNN is presented in the following theorem.

Theorem 1: *A given case C_2 is dominated by another case C_1 with respect to a query Q if and only if:*

$$\text{sim}(C_2, Q) + \sum_{a \in A - A_Q} w_a(1 - \text{sim}_a(C_1, C_2)) < \text{sim}(C_1, Q).$$

Proof: See Appendix A.

3.3. Attribute selection strategy

The attribute selected by iNN on each cycle of the recommendation process is the one that is most useful for confirming the case selected as the *target* case. The target case is first selected at random from the cases that are maximally similar to an initial query entered by the

user, and continually revised as the query is extended. No change is needed as long as the target case remains one of the cases that are maximally similar to the current query.

As Figure 2 illustrates, attribute selection in iNN aims to maximise the number of cases dominated by the target case. The cases currently dominated by the target case are shown in the lower half of the diagram. As indicated by the dashed arrows, there may be many dominance relationships with respect to the current query, but iNN considers only cases that are dominated by the target case. For each of the remaining attributes, it uses the dominance criterion from Theorem 1 to determine the number of cases that will be dominated by the target case if the preferred value of the attribute is the same as in the target case. It then selects the attribute that maximises the number of cases potentially dominated by the target case. If two or more attributes are equally promising according to this criterion, iNN uses the importance weights assigned to the case attributes as a secondary selection criterion. That is, it chooses the most important of the equally promising attributes.

3.4. Terminating the recommendation dialogue

As we have shown in previous work, naïve approaches to termination of recommendation dialogues such as stopping when the similarity of any case reaches a predefined threshold cannot guarantee that a better solution will not be found if the dialogue is allowed to continue (McSherry, 2003a). In fact, the only way to ensure that a more similar case (or another equally similar case) will not be found is to insist




Figure 2. Attribute selection in iNN aims to maximise the number of cases dominated by the target case.

that the recommendation dialogue is terminated only when the current query Q is such that:

$$r(Q^*) = r(Q) \text{ for all possible extensions } Q^* \text{ of } Q.$$

That is, the recommendation dialogue can be *safely* terminated only when it is certain that the recommendation will be the same no matter how the user chooses to extend her query.

It may seem at first sight that testing the above condition for safe termination of the recommendation dialogue may require an exhaustive search over all possible extensions of the current query. However, McSherry (2003a) shows that it can be tested without relying on exhaustive search and with a computational cost that increases only linearly with the size of the case library. The criteria used in iNN to recognise when the recommendation dialogue can be safely terminated are stated in the following theorem.

Theorem 2: *The recommendation dialogue in iNN can be safely terminated if and only if the following conditions hold:*

1. *any case that equals the similarity of the target case to the current query has the same values as the target case for all remaining attributes,*
2. *all cases that are less similar than the target case are dominated by the target case.*

Proof: See Appendix A.

Although expressed in terms of the target case in iNN's goal-driven approach to attribute selection, the criteria identified in Theorem 2 are equivalent to the criteria we have shown to be essential to ensure that the recommendation dialogue can be safely terminated in *any* approach to attribute selection (McSherry 2003a).

3.5. Recommendation efficiency

In previous work, we evaluated iNN in comparison with CCBP-PR algorithms based on a variety of different attribute-selection strategies (McSherry, 2003a). The performance measure of interest was recommendation *efficiency* as measured by the average number of questions the user is asked before the final recommendation is made. The algorithms compared differed only in their attribute-selection strategies, with termination of the recommendation dialogue based

on the criteria we have shown to be essential to ensure that the recommendation will remain unchanged in any approach to attribute selection.

Our evaluation was based on the Travel case library (www.ai-cbr.org), a standard benchmark containing more than 1,000 cases, and the PC case library (McGinty and Smyth, 2002). The results showed iNN to be more effective in reducing dialogue length than any of the other attribute-selection strategies. Its performance on the PC case library was close to optimal, reducing the number of questions asked by up to 63% and by 35% on average relative to a full-length query. It also gave the best performance on the Travel case library, reducing dialogue length by up to 63% and by 25% on average.

4. Explanation in Top Case

We now present an approach to explanation of the recommendation process in iNN in which explanations are automatically generated with no requirement for domain knowledge other than the similarity knowledge and cases already available to the system. We demonstrate the approach in a mixed-initiative recommender system called *Top Case* which can explain the relevance of any question the user is asked in strategic terms, recognise when the dialogue can be safely terminated, and justify its recommendations on the grounds that any remaining attributes cannot affect the outcome. An example recommendation dialogue based on a well-known case library in the travel domain is used to illustrate the approach.

4.1. *Explanation engineering*

An initial query entered by the user is incrementally *extended* in Top Case by asking the user to specify preferred values for attributes not mentioned in her initial query. On each recommendation cycle, the user is asked for the preferred value of the most useful attribute for confirming the target case and shown the competing cases that are now most similar to her query. The user can terminate the recommendation dialogue at any stage by selecting one of the cases she is shown as the product she prefers. Otherwise, query elicitation continues until Top Case has determined that its recommendation will be the same no matter how the user chooses to extend her query. At this point, the dialogue is terminated and the user is informed that the target case

has been confirmed as the recommended case. Any case that equals the current similarity of the target case at this stage is also recommended.

As well as being able to explain the relevance of any question it asks the user, Top Case can explain why it is recommending a particular case. Both types of explanation are dynamically generated at run time using explanation *templates*. The template used by Top Case to explain the relevance of a question the user is asked depends on whether or not the question has the potential to confirm the target case as the recommended case. Usually in the early stages of query elicitation, the selected attribute can at best be expected to increase the number of cases dominated by the target case rather than confirm the target case outright. The explanation template used when the target case cannot be confirmed in a single step is:

Because if $a = v$ this will increase the similarity of Case X from S_1 to S_2 {and eliminate N cases [including Cases X_1, X_2, \dots, X_r]}

where:

- a is the attribute whose preferred value the user is asked to specify,
- v is the value of a in the target case,
- Case X is the target case,
- S_1 is the similarity of the target case to the current query,
- S_2 is the similarity of the target case that will result if the preferred value of a is v ,
- N is the number of cases that will be eliminated if the preferred value of a is v ,
- Cases X_1, X_2, \dots, X_r are cases that the user was shown on the previous recommendation cycle that will be eliminated if the preferred value of a is v .

The section of the template enclosed in curly brackets is used only if one or more cases will be eliminated if the preferred value of a is v , which may not be the case in the early stages of query elicitation. The section enclosed in square brackets is used only if one or more of the cases that the user was shown in the previous recommendation cycle will be eliminated if the preferred value of a is v .

The template used by Top Case to explain the relevance of a question that does have the potential to confirm the target case in a single step is:

Because if $a = v$ this will confirm Case X as the recommended case

where a , v , and Case X are as defined for the previous template.

The template used by Top Case to explain why it is recommending a particular case depends on whether that case exactly matches the user's query and whether preferred values have been elicited for all case attributes. In the example template below, Case *X* is the recommended case, *attributes-1* are attributes in which it differs from the user's query, and *attributes-2* are attributes for which preferred values have not been elicited because they cannot affect the recommendation.

Case *X* differs from your query only in *attributes-1* and is the best case no matter what *attributes-2* you prefer

As in First Case (McSherry, 2003b), the aim of highlighting any compromises associated with the recommended case is to help the user decide whether or not to accept the system's recommendation. Also informing the user that her preferences with respect to any remaining attributes cannot affect the solution may help to increase her *confidence* in the recommendation. How effective the explanations provided by Top Case are in achieving these objectives is an important issue to be addressed by further research.

4.2. Example case library

The *Travel* case library (www.ai-cbr.org) that we use to illustrate our approach to explanation in Top Case is a standard benchmark containing the descriptions of over 1,000 holidays. Attributes in the case library and importance weights assigned to them in Top Case are price (8), month (7), location (6), persons (5), duration (4), type (3), accommodation (2), and transport (1).

Attributes price, persons, and duration are treated in Top Case as *nearer-is-better* (NIB) attributes. The remaining attributes are treated as nominal attributes. As usual in practice, we define the similarity of two given values x and y of a NIB attribute a to be:

$$\text{sim}_a(x, y) = 1 - \frac{|x - y|}{\max(a) - \min(a)},$$

where, for example, $\max(a)$ is the maximum value of a in the case library. Our similarity measure for each of the nominal attributes simply assigns a score of 1 to a given case if its value for the attribute is the same as the preferred value and 0 if it is not the same.

Figure 3 shows an example dialogue based on the *Travel* case library. Minor editing in the interest of clarity includes the format of the user's initial query. The current similarity of each retrieved case,

Top Case: Please enter your query
User: type = wandering, month = aug
Top Case: The target case is:
 Case 510: 609, aug, tyrol, 5, 7, wandering, flat, car (0.28)
 Other competing cases (972) include:
 Case 702: 1698, aug, giant_mountains, 3, 14, wandering, star2, car (0.28)
 Case 574: 529, aug, alps, 4, 7, wandering, flat, car (0.28)
 What is the preferred location?
User: **why**
Top Case: Because if location = tyrol this will increase the similarity of Case 510 from 0.28 to 0.44 and eliminate 866 cases, including Case 574
 What is the preferred location?
User: alps
Top Case: The target case is:
 Case 574: 529, aug, alps, 4, 7, wandering, flat, car (0.44)
 Other competing cases (82) include:
 Case 570: 799, aug, alps, 4, 7, wandering, flat, car (0.44)
 Case 586: 1958, aug, alps, 6, 14, recreation, flat, car (0.36)
 What is the preferred price?
User: **why**
Top Case: Because if price = 529 this will increase the similarity of Case 574 from 0.44 to 0.67 and eliminate 72 cases, including Case 570
 What is the preferred price?
User: 500
Top Case: The target case is:
 Case 574: 529, aug, alps, 4, 7, wandering, flat, car (0.67)
 Other competing cases (10) include:
 Case 74: 814, aug, alps, 6, 7, recreation, flat, train (0.57)
 Case 586: 1958, aug, alps, 6, 14, recreation, flat, car (0.54)
 What is the preferred number of persons?
User: **why**
Top Case: Because if number of persons = 4 this will confirm Case 574 as the recommended case
 What is the preferred number of persons?
User: 4
Top Case: The recommended case is Case 574 (0.8)
User: **why**
Top Case: Case 574 differs from your query only in price and is the best case no matter what transport, duration, or accommodation you prefer

Figure 3. Example recommendation dialogue based on the travel case library.

normalised by the sum of all the importance weights, is shown in brackets. In response to her initial query, the user is shown Case 510, the case initially selected by Top Case as the target case, and the two most similar of the 972 cases that are not currently dominated by the target case. Having taken the initiative, Top Case now selects location

as the most useful attribute for confirming the target case and asks the user what location she prefers.

When asked to explain the relevance of location, Top Case points out that if the preferred location is *Tyrol*, this will increase the similarity of the target case from 0.28 to 0.44 and eliminate 866 of the 972 competing cases. When the user chooses *Alps* instead as the preferred location, the target case changes to Case 574, but now there are only 82 competing cases. The user's answers to the next two questions are enough for Top Case to confirm Case 574 as the recommended case. When asked to explain its recommendation, Top Case points out that the recommended case differs from the user's query only in price and that her preferences with respect to the remaining attributes cannot affect the recommendation.

4.3. Discussion

A known limitation of similarity-based retrieval is that the most similar case may not be the one that is most acceptable to the user (e.g. McSherry, 2003b). It must also be recognised that the case recommended by Top Case may not be acceptable to the user even though it is guaranteed to remain the most similar case no matter how she chooses to extend her query. In future research we plan to investigate an approach to addressing this issue in which the dialogue is allowed to continue beyond the initial recommendation so that the user can extend or revise her query to include one or more constraints that must be satisfied.

It is worth noting that all three cases presented by Top Case in response to the user's initial query have the same values for holiday type and month; in fact there are 19 other cases that exactly match the user's initial query. Inseparability of competing cases is a common problem associated with incomplete queries (McSherry, 2002c) that highlights the importance of Top Case having the ability to take the initiative to help users discriminate between alternatives that are equally good in terms of their initial requirements.

5. Related Work

Allowing the user to enter an initial query to be incrementally extended is a feature that Top Case shares with mixed-initiative CCBR tools for fault diagnosis such as NaCoDAE (Aha et al.,

2001). Most of the dialogue features associated with mixed-initiative interaction in CBR (Aha et al., 2001; McSherry, 2001a, 2002a) are supported in Top Case, though not all are shown in the example dialogue. At any stage of the recommendation dialogue, for example, the user can specify a preferred value for an attribute other than the one considered most useful by Top Case. The user can also indicate her *indifference* to an attribute for which she is asked to specify a preferred value.

Sørmo et al. (2005, this issue) distinguish between different types of explanation in CBR according to the *goals* they support, such as explaining why the proposed solution is a good solution (*Justification*), explaining how the system reached the solution (*Transparency*), or explaining why a question is relevant (*Relevance*). The explanations provided by Top Case when asked why it is recommending a particular case can be seen to address the *Justification* goal. As well as addressing the *Relevance* goal, the explanations it provides when asked to explain the relevance of the questions it asks may also contribute to the *Transparency* goal of increasing the user's understanding of how the solution was obtained.

CBR Strategist (McSherry, 2001a) is a CCBR tool for fault diagnosis in which attribute selection is based on the reasoning strategies used by doctors, such as confirming a target diagnosis or eliminating a competing diagnosis (Elstein et al., 1978; Kassirer and Kopelman, 1991). As in iNN, an important benefit of CBR Strategist's *goal-driven* approach to attribute selection is that the relevance of any question the user is asked can be explained in terms of the purpose for which it was selected. Driven by an algorithm for strategic induction of decision trees (McSherry, 1999), CBR Strategist is best suited to diagnosis and classification tasks in which the number of outcome classes is small. This is not the case in product recommendation, where it is typical for each outcome class (a unique product or service) to be represented by a single case (McSherry, 2001b).

6. Conclusions

Following a brief discussion of existing approaches to explanation in recommender systems and lessons learned from this research, we have focused in this paper on the benefits of iNN, a CBR approach to product recommendation, in terms of making the recommendation process more transparent to users. We have presented a detailed account of

the recommendation process in iNN and how it combines an effective approach to reducing the length of recommendation dialogues with a mechanism for ensuring that the dialogue is terminated only when it is certain that the recommendation will be the same no matter how the user chooses to extend her query. We have also presented a novel approach to explanation of the recommendation process in which there is no requirement for domain knowledge other than the similarity knowledge and cases already available to the system.

We have demonstrated our approach in a mixed-initiative recommender system called *Top Case* which can explain the relevance of any question the user is asked in terms of its strategy of eliminating competing cases and ultimately confirming the target case as the recommended case. *Top Case* can also justify its recommendations on the grounds that any un-elicited preferences of the user cannot affect the outcome. In future research we plan to investigate the potential impact of the system's explanation capabilities and ability to support mixed-initiative interaction on the effectiveness of the recommendation process.

Appendix A. Theorems 1 and 2.

Lemma 1: *For any cases C_1, C_2 , query Q , and $a \in A$: $\text{sim}_a(C_2, Q) \leq \text{sim}_a(C_1, Q) + 1 - \text{sim}_a(C_1, C_2)$.*

Proof: By the triangle inequality, $1 - \text{sim}_a(C_1, Q) \leq 1 - \text{sim}_a(C_1, C_2) + 1 - \text{sim}_a(C_2, Q)$. The required inequality easily follows.

Theorem 1: *A given case C_2 is dominated by another case C_1 with respect to a query Q if and only if:*

$$\text{sim}(C_2, Q) + \sum_{a \in A - A_Q} w_a(1 - \text{sim}_a(C_1, C_2)) < \text{sim}(C_1, Q).$$

Proof: If the latter condition holds, then it must also be true that:

$$\text{sim}(C_2, Q) + \sum_{a \in A_{Q^*} - A_Q} w_a(1 - \text{sim}_a(C_1, C_2)) < \text{sim}(C_1, Q)$$

for any extension Q^* of Q . It follows from Lemma 1 that for any extension Q^* of Q :

$$\text{sim}(C_2, Q^*) = \text{sim}(C_2, Q) + \sum_{a \in A_{Q^*} - A_Q} w_a \text{sim}_a(C_2, Q^*)$$

$$\begin{aligned}
&\leq \text{sim}(C_2, Q) + \sum_{a \in A_{Q^*} - A_Q} w_a (\text{sim}_a(C_1, Q^*) + 1 - \text{sim}_a(C_1, C_2)) \\
&< \text{sim}(C_1, Q) + \sum_{a \in A_{Q^*} - A_Q} w_a \text{sim}_a(C_1, Q^*) = \text{sim}(C_1, Q^*).
\end{aligned}$$

So C_2 is dominated by C_1 as required. It remains to show that if:

$$\text{sim}(C_1, Q) \leq \text{sim}(C_2, Q) + \sum_{a \in A - A_Q} w_a (1 - \text{sim}_a(C_1, C_2))$$

then C_2 is not dominated by C_1 . Let Q^* be the *complete* extension of Q such that $\pi_a(Q^*) = \pi_a(C_2)$ for all $a \in A - A_Q$. It can be seen that $\text{sim}_a(C_2, Q^*) = 1$ and $\text{sim}_a(C_1, Q^*) = \text{sim}_a(C_1, C_2)$ for all $a \in A - A_Q$, and so:

$$\begin{aligned}
\text{sim}(C_1, Q^*) &= \text{sim}(C_1, Q) + \sum_{a \in A - A_Q} w_a \text{sim}_a(C_1, Q^*) \\
&\leq \text{sim}(C_2, Q) + \sum_{a \in A - A_Q} w_a (\text{sim}_a(C_1, Q^*) + 1 - \text{sim}_a(C_1, C_2)) \\
&= \text{sim}(C_2, Q) + \sum_{a \in A - A_Q} w_a = \text{sim}(C_2, Q^*).
\end{aligned}$$

It follows as required that C_2 is not dominated by C_1 .

Theorem 2: *The recommendation dialogue in iNN can be safely terminated if and only if the following conditions hold:*

1. *any case that equals the similarity of the target case to the current query has the same values as the target case for all remaining attributes,*
2. *all cases that are less similar than the target case are dominated by the target case.*

Proof: By definition, the recommendation dialogue can be safely terminated if and only if the current query Q is such that $r(Q^*) = r(Q)$ for all possible extensions Q^* of Q . Also by definition, $C_t \in r(Q)$, where C_t is the case currently selected by iNN as the target case.

Suppose now that Conditions 1 and 2 hold, and let Q^* be any extension of Q . It is clear from Condition 1 that $\text{sim}(C, Q^*) = \text{sim}(C_t, Q^*)$ for any $C \in r(Q)$. On the other hand, it follows from Condition 2 that for any $C \notin r(Q)$, $\text{sim}(C, Q^*) < \text{sim}(C_t, Q^*)$ and so $C \notin$

$r(Q^*)$. We have established that $r(Q) = r(Q^*)$ for any extension Q^* of Q and so the dialogue can be safely terminated as required.

If Condition 1 is not satisfied, then there exists $C \in r(Q)$ and $a^\circ \in A - A_Q$ such that $\pi_{a^\circ}(C) \neq \pi_{a^\circ}(C_t)$. It follows that if Q^* is the *complete* extension of Q such that $\pi_a(Q^*) = \pi_a(C)$ for all $a \in A - A_Q$, then $\text{sim}_{a^\circ}(C_t, Q^*) < \text{sim}_{a^\circ}(C, Q^*) = 1$, and so $\text{sim}(C_t, Q^*) < \text{sim}(C, Q^*)$. Thus $C_t \notin r(Q^*)$. It follows that Condition 1 is a necessary condition for the dialogue to be safely terminated.

If Condition 2 is not satisfied, then there exists $C \notin r(Q)$ such that C is not dominated by C_t . As $\text{sim}(C_t, Q) > \text{sim}(C, Q)$, there must be an extension Q^* of Q such that $\text{sim}(C, Q^*) \geq \text{sim}(C_t, Q^*)$. It follows that $r(Q^*) \neq r(Q)$, so the dialogue cannot be safely terminated. Thus Condition 2 is also a necessary condition for the dialogue to be safely terminated.

References

- Aha, D. W., Breslow, L. A. & Muñoz-Avila, H. (2001). Conversational Case-Based Reasoning. *Applied Intelligence* **14**: 9–32.
- Armengol, E., Paludàries, A. & Plaza, E. (2001). Individual Prognosis of Diabetes Long-Term Risks: a CBR Approach. *Methods of Information in Medicine* **40**: 46–51.
- Burke, R. (2002). Interactive Critiquing for Catalog Navigation in E-Commerce. *Artificial Intelligence Review* **18**: 245–267.
- Doyle, M. & Cunningham, P. (2000). A Dynamic Approach to Reducing Dialog in On-Line Decision Guides. In Blanzieri, E. & Portinale, L. (eds.) *Advances in Case-Based Reasoning*, 49–60. Springer-Verlag: Berlin Heidelberg.
- Elstein, A. S., Schulman, L. A. & Sprafka, S.A. (1978). *Medical Problem Solving: an Analysis of Clinical Reasoning*. Harvard University Press: Cambridge, MA.
- Evans-Romaine, K. & Marling, C. (2003). Prescribing Exercise Regimens for Cardiac and Pulmonary Disease Patients with CBR. In McGinty, L. (ed.) *ICCB-03 Workshop Proceedings*, 45–52. Technical Report 4/2004, Department of Computer and Information Science, Norwegian University of Science and Technology.
- Gaasterland, T., Godfrey, P. & Minker, J. (1992). An Overview of Cooperative Answering. *Journal of Intelligent Information Systems* **1**: 123–157.
- Hammond, K. J., Burke, R. & Schmitt, K. (1996). A Case-Based Approach to Knowledge Navigation. In Leake, D.B. (ed.) *Case-Based Reasoning: Experiences, Lessons & Future Directions*, 125–136. AAAI Press: Menlo Park, CA.
- Herlocker, J. L., Konstan, J. A. & Riedl, J. (2000). Explaining Collaborative Filtering Recommendations. In *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work*, 241–250. ACM Press: New York, NY.
- Kassirer, J. P. & Kopelman, R. I. (1991). *Learning Clinical Reasoning*. Williams and Wilkins: Baltimore, MD.
- Kohlmaier, A., Schmitt, S. & Bergmann, R. (2001). A Similarity-Based Approach to Attribute Selection in User-Adaptive Sales Dialogues. In Aha, D.W. & Watson, I.

- (eds) *Case-Based Reasoning Research and Development*, 306–320. Springer-Verlag: Berlin Heidelberg.
- McGinty, L. & Smyth, B. (2002). Comparison-Based Recommendation. In Craw, S. & Preece, A. (eds.) *Advances in Case-Based Reasoning*, 575–589. Springer-Verlag: Berlin Heidelberg.
- McSherry, D. (1999). Strategic Induction of Decision Trees. *Knowledge-Based Systems* **12**: 269–275.
- McSherry, D. (2001a). Interactive Case-Based Reasoning in Sequential Diagnosis. *Applied Intelligence* **14**: 65–76.
- McSherry, D. (2001b). Minimizing Dialog Length in Interactive Case-Based Reasoning. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*, 993–998. Morgan Kaufmann: San Francisco, CA.
- McSherry, D. (2002a). Mixed-Initiative Dialogue in Case-Based Reasoning. In Aha, D.W. (ed.) *Proceedings of the ECCBR-02 Workshop on Mixed-Initiative Case-Based Reasoning*, 1–8. Robert Gordon University: Aberdeen.
- McSherry, D. (2002b). Recommendation Engineering. In *Proceedings of the Fifteenth European Conference on Artificial Intelligence*, 86–90. IOS Press: Amsterdam.
- McSherry, D. (2002c). The Inseparability Problem in Interactive Case-based Reasoning. *Knowledge-Based Systems* **15**: 293–300.
- McSherry, D. (2003a). Increasing Dialogue Efficiency in Case-Based Reasoning Without Loss of Solution Quality. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, 121–126. Morgan Kaufmann: San Francisco, CA.
- McSherry, D. (2003b). Similarity and Compromise. In Ashley, K.D. & Bridge, D.G. (eds.) *Case-Based Reasoning Research and Development*, 291–305. Springer-Verlag: Berlin Heidelberg.
- McSherry, D. (2004). Incremental Relaxation of Unsuccessful Queries. In Funk, P. & González-Calero, P. (eds.) *Advances in Case-Based Reasoning*, 331–345. Springer-Verlag: Berlin Heidelberg.
- Reilly, J., McCarthy, K., McGinty, L. & Smyth, B. (2005). Explaining Compound Critiques. *Artificial Intelligence Review*. This Issue.
- Shimazu, H. (2002). ExpertClerk: A Conversational Case-Based Reasoning Tool for Developing Salesclerk Agents in E-Commerce Webshops. *Artificial Intelligence Review* **18**: 223–244.
- Sørmo, F. & Aamodt, A. (2002). Knowledge Communication and CBR. In González-Calero, P. (ed.) *Proceedings of the ECCBR-02 Workshop on Case-Based Reasoning for Education and Training*, 47–59. Robert Gordon University: Aberdeen, Scotland.
- Sørmo, F., Cassens, J. & Aamodt, A. (2005). Explanation in Case-Based Reasoning: Perspectives and Goals. *Artificial Intelligence Review*. This Issue.